Redpaper

Ilya Krutov
Raj Panda
Arthur Ban
Ruzhu Chen
Zaphiris Christidis
Bradley Elkin
Tony Pirraglia
Hari Reddy

Positioning IBM System x3750 M4 for High Performance Computing Workloads

The purpose of this IBM® Redpaper™ publication is to discuss why IBM System x3750 M4 server is a reasonable choice for the majority of high performance computing (HPC) workloads.

In this paper we analyze capabilities of x3750 M4 that are particularly well suited for HPC deployments, with proof points derived from the performance results obtained under various HPC workloads.

The following topics are covered:

Executive summary

High Performance Computing (HPC) workloads are compute-intensive modelling and simulation workloads applicable to various field areas such as high energy physics, weather forecasting, computer-aided engineering, genome sequencing, computational chemistry, and others. These workloads commonly deal with significant amounts of input data that needs to be processed in a most efficient way. Efficient means the optimal combination of factors such as processing time required, electrical energy used, space occupied, and cost incurred to build and manage the computing system itself. In other words, the HPC system should provide sufficient raw compute power (including processor, memory and I/O throughput) in a compact energy-efficient package at a reasonable cost, matching the workload size.

Depending on workload size, the computing systems can range from simple single-purpose servers to complex multinode multipurpose compute clusters employing high-speed node interconnects, such as InfiniBand.

The IBM System x3750 M4 server, built on the Intel Xeon processor E5-4600 product family, delivers an extensive set of HPC features to provide an optimal balance between performance, energy consumption, occupied space, and acquisition and management costs for a broad range of HPC workloads being either initially deployed or moved or consolidated from existing implementations. While IBM BladeCenter®-based and IBM iDataPlex®-based HPC cluster deplyments are good for large scale implementations, the x3750 M4 is a reasonable choice for small single- or multinode workgroup HPC clusters.

To explore HPC capabilites of the x3750 M4, a set of HPC benchmarks was run on it, and obtained results were compared with the results of other dual- and quad-socket IBM System x® servers (both current and previous generations). The following HPC benchmarks were run on x3750 M4 as part of our analysis:

- LINPACK
- Weather (WRF)
- Quantum Chromodynamics (MILC)
- Next Generation Sequencing (NGS)
- Computational Chemistry (GAUSSIAN)
- Computer Aided Engineering (FLUENT)

We focused on the performance scalability of x3750 M4 itself and in relation to other dual- and quad-socket systems, and observed the following:

- The floating point performance of 4-socket x3750 M4 with Intel Xeon processor E5-4650 is almost equivalent to the performance of two 2-socket dx360 M4 or x3550 M4 with Intel Xeon processor E5-2680.
- Single IBM System x3750 M4 can consolidate HPC workloads from up to five 2-socket dx360 M3 servers with Intel Xeon X5670 processors.
- IBM System x3750 M4 can provide more than three times improvement in technical computing performance compared to the previous generation of x3850 X5 with Intel Xeon X7560 processors.
- The x3750 M4 server can achieve almost linear performance scalability for HPC workloads going from eight to 16 to 32 cores.

Depending on workload size, x3750 M4 can potentially eliminate the need for multinode HPC workload deployments or significantly reduce the number of nodes required to support specific workloads, therefore providing savings on the external interconnect infrastructure and making the deployment, support and maintenance tasks easier, especially with workgroup HPC clusters.

Introduction

High Performance Computing (HPC) workloads are compute-intensive modelling and simulation workloads applicable to various field areas such as high energy physics, weather forecasting, computer-aided engineering, genome sequencing, computational chemistry, and others. These workloads commonly deal with significant amounts of input data that needs to be processed in a most efficient way. Efficient means the optimal combination of factors such as processing time required, electrical energy used, space occupied, and cost incurred to build and manage the computing system itself. In other words, the HPC system should provide sufficient raw compute power (including processor, memory and I/O throughput) in a compact energy-efficient package at reasonable cost, matching the workload size.

Depending on workload size, the computing systems can range from simple single-purpose servers to complex multinode multipurpose compute clusters employing high-speed node interconnects, such as InfiniBand.

The IBM System x3750 M4 server, built on the Intel Xeon processor E5-4600 product family, deliver an extensive set of HPC features to provide optimal balance between performance, energy consumption, occupied space, and acquisition and management costs, for a single-purpose single-system HPC workloads either being initially deployed or being moved or consolidated from the multinode implementations.

IBM System x3750 M4 overview

The IBM System x3750 M4 is a 4-socket server featuring a streamlined design, optimized for price and performance, with best-in-class flexibility and expandability. Models of the x3750 M4 are powered with the Intel Xeon processor E5-4600 product family, up to eight cores each, for an entry level 4-socket solution. The x3750 M4 maximizes storage density and provides flexible PCI and 10 Gb Ethernet networking options in a 2U form factor.

Figure 1 shows the IBM System x3750 M4.

Figure 1 The IBM System x3750 M4

The x3750 M4 has exceptional memory performance that is achieved by supporting three-RDIMM-per-channel configurations at speeds up to 25% faster than the Intel specification, while still maintaining world-class IBM reliability. LR-DIMM speeds are also 25% beyond the Intel specification for 1.35 V DIMMs, and this speed improves not only performance, but reduces overall system power at the same time.

The x3750 M4 offers a flexible, scalable design and a simple upgrade path to 16 hard-disk drives (HDDs) or 32 IBM eXFlash solid-state drives (SSDs), with up to eight PCIe 3.0 slots and up to 1.5 TB of memory. The flexible embedded Ethernet solution provides two standard Gigabit Ethernet ports onboard, along with a dedicated 10 GbE slot that allows for a choice of either two copper or two fiber optic connections. Comprehensive systems management tools with the next-generation Integrated Management Module II (IMM2) make it easy to deploy, integrate, service, and manage.

The IBM System x3750 M4 blends outstanding flexibility and expandability. The x3750 M4 2+2 socket design enables pay-as-you-grow processing with the new Intel Xeon processor E5-4600 product family and memory scalability to help lower cost and manage growth. The 5+3 PCIe socket design allows you to pay for PCIe capabilities as needed.

IBM System x3750 M4 positioning

IBM System x3750 M4 is introduced to fill the demand for application workloads that require 4-socket processor scalability but do not require advanced RAS and scalability features found in IBM eX5 enterprise servers. One of the major classes of applications that fits this definition is HPC workloads.

Table 1 compares key specifications of standard 2-socket x3550 M4, 4-socket x3750 M4 and up to 8-socket scalable x3850 X5 rack servers, as well as IBM BladeCenter HS23 and IBM iDataPlex dx360 M4 nodes.

Table 1 Specification comparison: HS23, dx360 M4, x3550 M4, x3750 M4, and x3850 X5

Component	HS23	dx360 M4	x3550 M4	x3750 M4	x3850 X5[a]
Form factor	Blade (14/9U rack)	iDataPlex node (2/2U rack)	1U rack	2U rack	4U rack
Processor family	Intel Xeon E5-2600	Intel Xeon E5-2600	Intel Xeon E5-2600	Intel Xeon E5-4600	Intel Xeon E7-8800
Max number of CPUs	2	2	2	4	4
DIMM slots	12	16	24	48	64 or 96[b]
Max memory capacity	256 GB	512 GB	768 GB	1.5 TB	2 TB or 3 TB[b]
Memory speed at max capacity	1600 MHz	1333 MHz	1066 MHz	1333 MHz	1066 MHz
2.5-in. HDDs/SSDs	2	2	8	16	8
1.8-inch SSDs (IBM eXFlash)	0	4	0	32	16
Max eXFlash IOPS	0	120,000	0	960,000	480,000
Max PCIe slots bandwidth	23.6 GBps	31.5 GBps	47.3 GBps	63 GBps	32 GBps
RAS	Standard	Standard	Standard	Standard	Advanced
Vertical scalability	None	None	None	2 + 2 socket plus memory	Scalable to 8-way, additional memory scalability with MAX5

a. x3850 X5 numbers are shown for the 4-way server. For 8-way system, they need to be multiplied by 2.
b. 96 DIMM slots and 3 TB of memory with an optional MAX5 memory expansion unit

From the specification point of view, x3750 M4 fits right in the middle between x3550 M4 and x3850 X5, providing affordable performance scalability along with standard RAS features. At the same time, up to 25% higher memory speed, twice as much IOPS and 33 to 50% higher I/O bandwidth in a dense 2U package makes the x3750 M4 a perfect choice for HPC workloads.

In large scale computing environments, BladeCenter HS23-based deployment provides the highest level of component integration while iDataPlex dx360 M4-based deployment provides the smallest footprint for HPC workloads. However, in small (also called *workgroup*) HPC deployments, IBM System x3750 M4 can provide better value than blade- or traditional 2-socket rack-based environments by reducing the number of servers required to deploy the specific workload, and therefore reducing the number of infrastructure components such as network adapters, interconnect links and switch ports, which in turn helps to lower capital expenses as well as operational and management costs while keeping required performance levels.

As an example, let us assume that you need to deploy a workgroup HPC cluster capable of 4 TFLOPS, you plan to use InfiniBand as a cluster interconnect, and you calculate theoretical LINPACK performance using the following formula:

```
GFLOPS = <number of cores> x <core frequency, GHz> x <number of floating point
operations per cycle>
```

Then you calculate the required number of nodes, adapters, cables and external network ports for each of the following approaches: BladeCenter HS23, iDataPlex dx360 M4, and x3750 M4. The results of our calculations are presented in Table 2.

Table 2 Workgroup HPC cluster comparison: HS23, dx360 M4, and x3750 M4

Component	HS23	dx360 M4	x3750 M4
Processor	E5-2680, 2.7 GHz	E5-2680, 2.7 GHz	E5-4650, 2.7 GHz
Number of processors per node	2	2	4
Number of cores per node	16	16	32
Calculated node performance	345.6 GFLOPS	345.6 GFLOPS	691.2 GFLOPS
Calculated cluster performance	4.1 TFLOPS	4.1 TFLOPS	4.1 TFLOPS
Number of nodes	12	12	6
Number of InfiniBand adapters	12	12	6
Number of cables	0	12	6
Number of InfiniBand ports	12	12	6

Clearly, twice fewer compute nodes, InfiniBand adapters and external switch ports helps to lower acquisition costs as well as deployment, maintenance, and support costs.

HPC benchmarks

High Performance Computing (HPC) is about using fast computers to solve complex problems in design, or simulate physical phenomena for understanding the behavior of such phenomena. HPC techniques are being used in many fields today including Computer-Aided Engineering (CAE), high energy physics, weather and climate modeling, to name a few. Some metric or benchmark is typically used to assess the compute capability of HPC systems. Since HPC systems are used in diverse fields, there is no one benchmark that can be used to assess how capable a given system is across all areas of use.

Therefore, IBM Performance Lab conducted series of tests utilizing different HPC benchmarks to highlight performance and scalability features of the IBM System x3750 M4. The following benchmarks were run:

- LINPACK
- Weather (WRF)
- Quantum Chromodynamics (MILC)
- Next Generation Sequencing (NGS)
- Computational Chemistry (GAUSSIAN)
- Computer Aided Engineering (FLUENT)

The benchmark results were evaluated for scalability, that is, how performance scales when increasing the number of compute cores from 8 to 16 to 32. Some tests also provided qualitative comparisons with current 2-way systems and previous generations of 2-way and 4-way systems.

The tested x3750 M4 server had the configuration shown in Table 3.

Table 3 x3750 M4 tested configuration

Specification	Description
Processor model	Intel Xeon processor E5-4650 @ 2.7GHz
Number of processors	4
Memory installed	256 GB (32x8 GB) @ 1600 MHz (DDR3)

Other tested systems used for reference purposes had the same number of processor cores to highlight the differences in scale up versus scale out approaches, as well as performance improvement over the previous generation.

Overall, x3750 M4 demonstrated clear advantages over its predecessors in both performance and scalability metrics. The following subsections discuss this in details.

LINPACK

There are a few benchmarks that HPC users do look at to get a "first impression" of a system, and one of them is LINPACK.

LINPACK is probably the best known HPC benchmark. It measures how fast a computer can solve dense systems of linear equations and is an indicator of the floating point capability of a computing system.

We then ran LINPACK on the IBM x3750 M4 system to get a performance of 661 GFLOPS when run on all the 32 cores in the system. The result is available at:

http://public.dhe.ibm.com/eserver/benchmarks/news/newsblurb_x3750M4_LINPACK_20120928.pdf

We also extracted LINPACK performance data for dx360 M4, dx360 M3 and x3850 X5 from the LINPACK performance report available at:

http://www.netlib.org/benchmark/performance.ps

For multinode clusters, we calculated the average LINPACK result per node. The comparison of these results is shown in Figure 2 (higher is better).

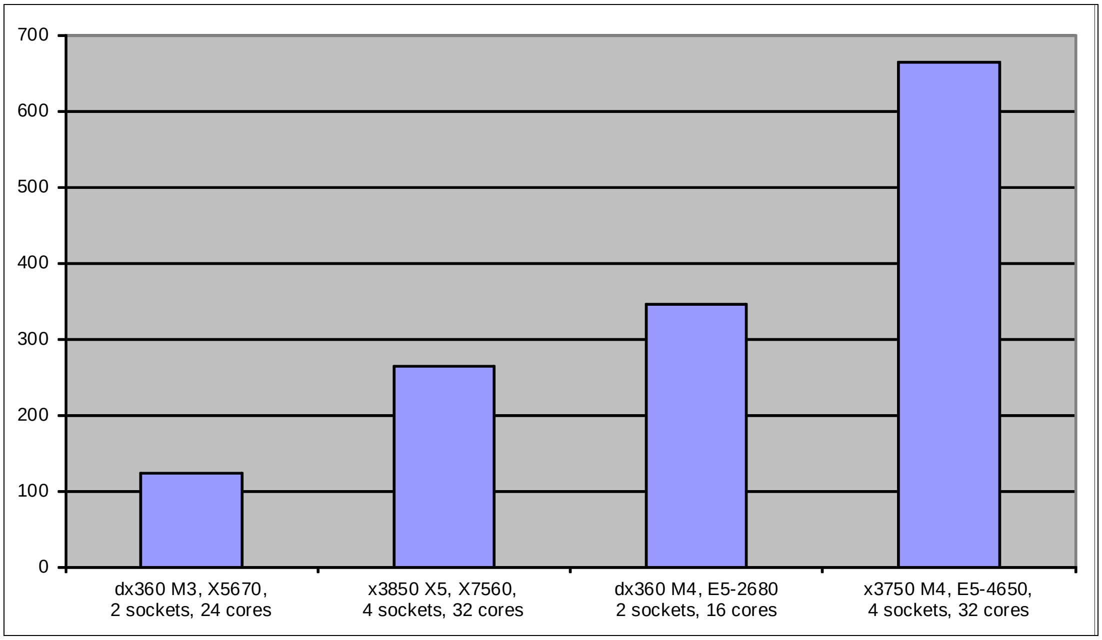

Figure 2 LINPACK performance evaluation: dx360 M3, x3850 X5, dx360 M4, and x3750 M4

We can conclude that one 4-way x3750 M4 (32 cores total) built on Intel Xeon processor E5-4650 can have almost the same performance as two 2-way dx360 M4s (32 cores total) running Intel Xeon processor E5-2680 without the need to build a two-node cluster, and it can have the performance comparable to five clustered dx360 M3 nodes that have a total of 60 cores, based on the Intel Xeon X5670 processor. Also, x3750 M4 performs more than 2.5 times better than 4-socket x3850 X5 running Intel Xeon X7560 processors.

We used x3750 M4's BIOS settings shown in Table 4 as well as a matrix size of N=160,000 to get the optimal performance. It should be noted that the large amount of memory of 256 GB available in the system enabled us to run such a large problem size with LINPACK.

Table 4 BIOS settings for optimal performance of LINPACK

Parameter	**Value**
Intel Hyper-Threading Technology	Disabled
Prefetcher	Enabled
NUMA	Enabled
Intel Turbo Boost Technology	Enabled

Weather (WRF)

The Weather Research and Forecasting (WRF) Model is a next-generation mesoscale numerical weather prediction system designed to serve both operational forecasting and atmospheric research needs. It features multiple dynamic cores, a 3-dimensional variational (3DVAR) data assimilation system, and a software architecture allowing for computational parallelism and system extensibility. WRF is suitable for a broad spectrum of applications across scales ranging from meters to thousands of kilometers.

WRF gives researchers the ability to conduct simulations reflecting either real data or idealized configurations. WRF provides operational forecasting a model that is flexible and efficient computationally, while offering the advances in physics, numerics, and data assimilation contributed by the research community.

WRF has a rapidly growing community of users, and workshops and tutorials are held each year at NCAR. WRF is currently in operational use at NCEP. This site provides information about the WRF effort and its organization, references to projects and forecasting involving WRF, and links to the WRF users' page, real-time applications, and WRF-related events.

The WRF code was installed on several IBM systems, and the WRF 12KM Continental US (CONUS) benchmark data was downloaded from the NCAR site at:

http://www.mmm.ucar.edu/wrf/WG2/benchv3

In all tests, Gigaflops per second (GFLOPS) were derived by averaging the computational timings for each time step. Thus the average time per time step is expressed as the "mean" value. Simulation speed was defined as the model time step, 72 seconds, divided by average time per time step. GFLOPS was defined as the simulation speed times 0.418 for this case based on a measured operation count of 30.1 billion floating point operations per second (or simply the operation count divided by the average time per time step).

WRF was tested on a 32-core x3850 X5 with Intel Xeon X7560 2.27 GHz processors and on a 32-core x3750 M4 built on Intel Xeon processor E5-4650 rated at 2.7 GHz and running in Turbo mode. In addition, and to provide a direct comparison with a 32-core previous generation system, WRF runs were conducted on an IBM iDataPlex dx360 M3 system consisting of three Intel Xeon X5670 (2.93 GHz) processor-based nodes interconnected with QDR Infiniband. This system consisted of two sockets (12 cores) per node.

To take into account the architectural differences of the new Intel Xeon processors that support Advanced Vector eXtensions (AVX), WRF was also compiled and executed with the -xAVX flag enabled, replacing the -xSSE4.2 flag, in the compiler options. The -xAVX flag produced appropriate vector binary code for the Intel Xeon processor E5 family, and proved to increase the overall performance by a significant factor.

Figure 3 shows the graphical representation of relative performance results obtained in GFLOPS (higher is better) on all systems and for 8, 16, and 32 MPI tasks (8-core x3850 X5 result is the baseline).

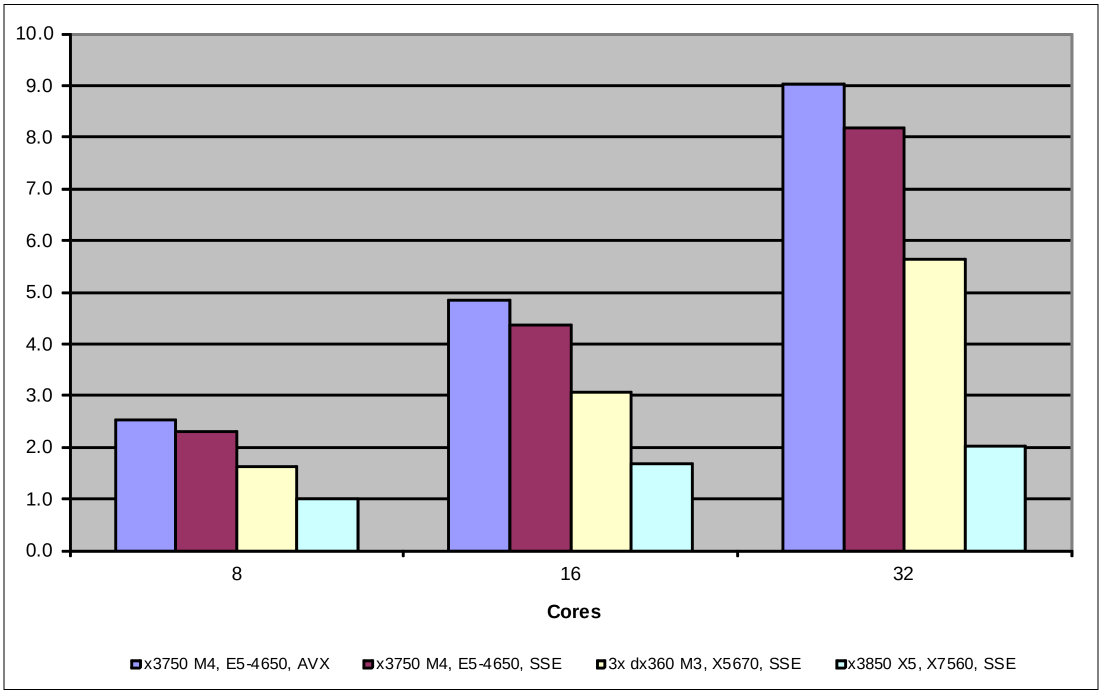

Figure 3 WRF benchmark results

Observing the results, we see an almost linear performance increase from 8 to 32 MPI tasks for x3750 M4. In addition, x3750 M4 can perform up to 1.6 times faster than a three-node cluster consisting of Intel Xeon X5670 processor-based compute nodes, and it can perform up to 4.5 times faster than x3850 X5 running Intel Xeon X7560 processors. These results proved that x3750 M4 can be a good choice when deploying new or consolidating existing weather simulation workloads.

Quantum Chromodynamics (MILC)

High Energy Physics (also known as Particle Physics) explores the properties and relationships of the building blocks of what we commonly refer to as matter and energy. In the current state-of-the-art understanding, particles are associated with quantum fields. The fundamental particles (from which all others are derived) comprise the Standard Model of particle physics. Describing (by theory) and confirming (by experiment) the properties of the Standard Model and physics beyond the Standard Model is the main focus of the high energy physics community. Theoreticians work closely with experimentalists to check the predictions of their models, and are distributed among many research institutions around the world. Some of the sites focusing on high energy physics, and another community applies similar theoretical tools (along with other models such as general relativity) to understand the astrophysics of phenomena observed by astronomers, like the interior of neutron stars.

Experimental physicists need large amounts of computing power to process the data they generate (a textbook case of "big data" in physics), but the main users of compute cycles are theorists. Quantum field theories like quantum chromodynamics (QCD) are extremely compute-intensive. Lattice QCD is the common approach used to solve the QCD theory of subatomic particles (the strong interactions of quarks and gluons). Unlike analytical and perturbation techniques, it can be applied to low-energy QCD. It uses Monte-Carlo and molecular dynamics algorithms to sample gauge configurations, which requires supercomputing power. Some of the programs and libraries supporting Lattice QCD simulations are MILC, CPS, Chroma, and FermiQCD.

MILC (MIMD Lattice Computation) works on four dimensional grids, integrating a differential equation of motion over thousands of time steps. Each step inverts a large, sparse matrix using the conjugate gradient method, but because the linear system is ill-conditioned, it requires many conjugate gradient (CG) iterations to converge to a solution. Executing the CG steps dominates the compute time. On a massively parallel system, the lattice is divided into sublattices distributed across each separate node. Within the sublattice used by a single processor, the four-dimensional nature of the grid requires data to be gathered from widely separated grid and link locations, which requires a significant amount of memory traffic, and when requesting data from off-node, a significant amount of network traffic between system nodes. One key concern is in the partitioning of sublattices across system nodes; the larger the sublattice, the higher the ratio of data used locally on the node versus data provided from the network. But larger sublattices can reduce single-core performance when they reduce cache efficiency. So a balanced system that can support larger sublattices locally is important to supporting a system that can scale.

The standard MILC benchmark from SPEC MPI2007 vsn 69 (SPEC MPI2007 vsn 1.1) was run on an x3750 M4 node and a previous-generation (x3850 X5) node. On both systems, MILC was built with the Intel Compiler suite, and x3850 X5 uses SSE while x3750 M4 utilizes AVX.

The benchmarks were managed with Intel MPI—for an 8-way run, tasks were started and each bound to a specific processor core. All benchmarks were run to minimize the number of sockets used. Both platforms have eight cores available per processor socket. So, for example, a 16-way run saturated two sockets (with two sockets left idle).

Figure 4 shows relative benchmark result comparisons for 8, 16, and 32 tasks (lower is better).

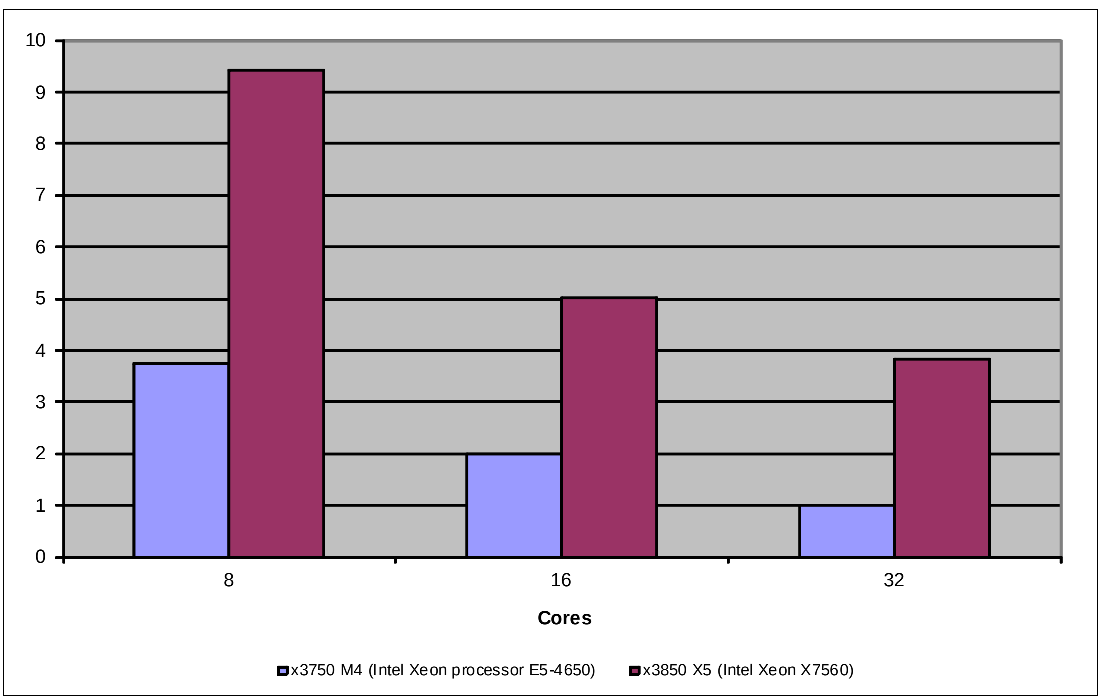

Figure 4 MILC benchmark results

The performance advantages of the x3750 M4 system are clear. First, for smaller task counts (8, 16), the x3750 is 2.5 times faster than previous generation 4-socket platforms. But more importantly, the fully saturated x3750 node does not hit a performance bottleneck when running 32 tasks. It reaches nearly the same scaling efficiency as measured for 16 tasks. This scaling behavior is critical for supporting massively parallel applications like MILC across large numbers of nodes.

Next Generation Sequencing (NGS)

The evolutional Next Generation Sequencing (NGS) technology prompts researchers and scientists to seek highly optimized computing solutions to handle immense volumes of NGS data. These especially include fast and reliable sequence assembly and mapping tools. Accurate sequence assembly, mapping and alignment of NGS sequence data (that is, reads) are time-consuming and require high performance computations and storage technology as well as efficient software tools to work together. To meet these requirements, a high performance computing system, as well as highly optimized sequence assembly and mapping software are necessary to handle the data workload.

CLC Assembly Cell is a licensed product developed by CLC bio in Denmark. It is a well known sequence assembly and mapping application used by more than 1700 universities and research agencies worldwide. It is implemented with Intel SIMD (single instruction, multiple data) to parallelize and accelerate sequence assembly and mapping algorithms, making it faster than many other widely used sequence assembly and mapping software. IBM, offering customers IT solutions and service, teamed with CLC bio to create optimized NGS solutions for speeding up sequence assembly and mapping on our system and storage technology.

CLC Assembly Cell is a package of command line tools that take sequence data from various formats for assembly and mapping. These include programs for de novo assembly, reference mapping (or alignment) and many other tools. De novo assembly enables billions of short reads (that is, short sequences) to be assembled quickly without reference sequences. The program is parallelized with POSIX threads (pthreads) and can run on a single node with multiple processor cores. Reference assembly, on the other hand, enables billions of short reads to be aligned or mapped to a known sequence. The reference mapping is also parallelized with pthreads and the job can be divided into smaller jobs and run on multiple nodes with multiple cores. The mapping outputs on each node are then merged.

The program requires high throughput storage and file systems such IBM GPFS™, x86 processors and moderate memory usage. This makes x3750 M4 a perfect fit.

CLC Assembly Cell 3.3 and 4.0 beta were used for benchmark test on the latest IBM system x3750 M4. The benchmark test case is 37x coverage of human genome sequence reads generated by Illumina.

The 2.7 GHz, 32-core x3750 M4 installed with 256 GB DDR3 memory was booted with turbo mode and hyper threading enabled. The benchmark tests were executed in an ext3 file system created on a local SSD drive. To better understand the performance on x3750 M4, the benchmarks were also executed on a 2.27 GHz, 32-core x3850 X5 system with 1.5 TB DDR3 memory and a 3.3 TB local xfs file system striped across multiple SAS disks.

Two options, paired-end (Figure 5 on page 12) and non paired-end (Figure 6 on page 13), were tested for de novo assembly, while only the paired-end option was used for reference mapping.

Using non paired-end information, the inputs are loaded and processed without pairing together. Figure 5 on page 12 shows a relative view of these results. The x3750 M4 outperformed x3850 X5 by around 5%. Using multiple threads from 16 to 64 on a 32 physical core system, the results showed more than a two-fold speed up, indicating that the hyperthreading helps improve the de novo assembly performance by approximately 20 to 25%.

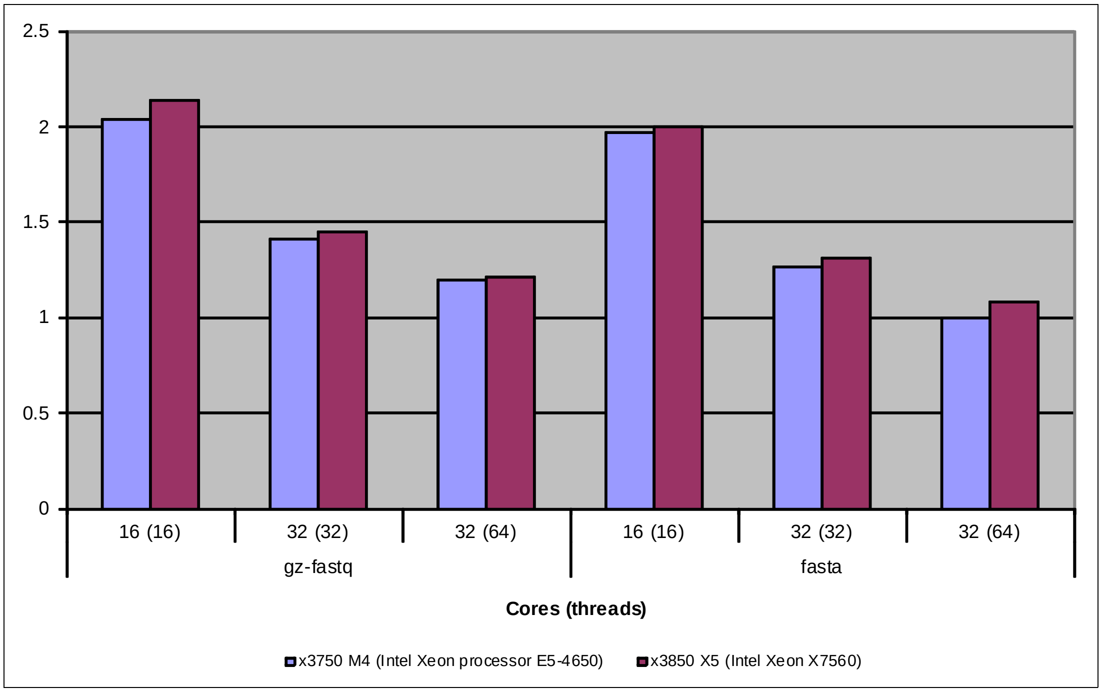

Figure 5 Non paired-end information benchmark results

The performance of sequence assembly on the two test systems is only marginally different. The reason may be that the CLC Assembly Cell program was built on the older version of the Intel system. It does not utilize the new features of Intel Xeon processor E5-4600 product family technology. In addition, the difference in I/O performance of file systems on the two systems may also affect the performance, since heavy I/O activity was seen during the processes.

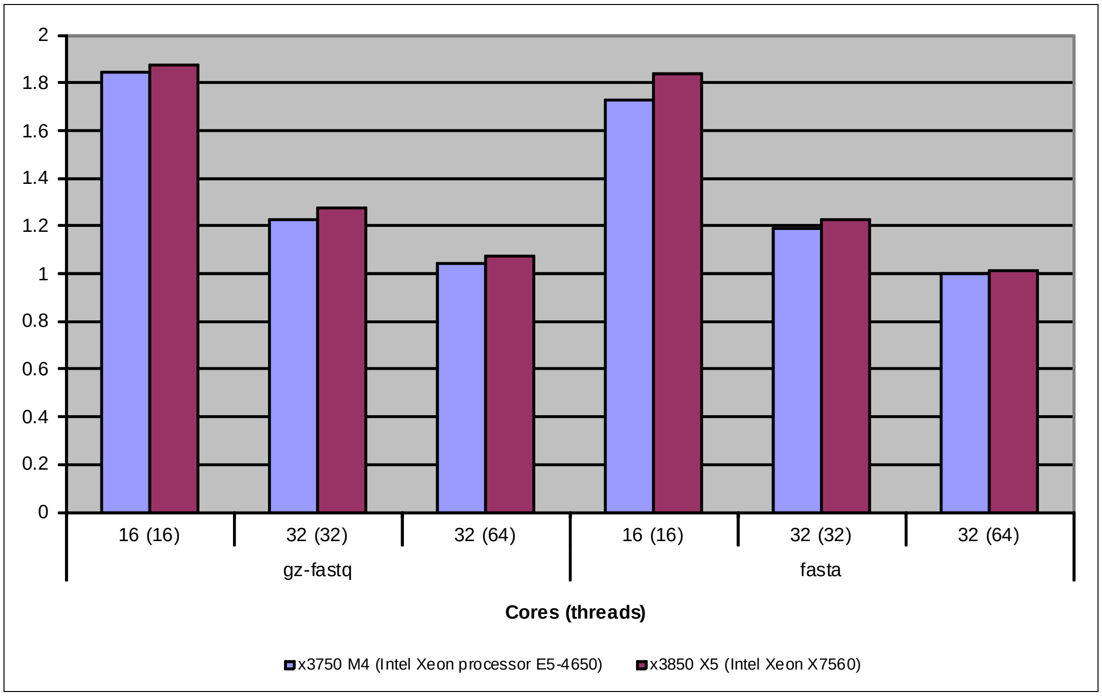

Figure 6 Paired-end information benchmark results

The benchmark with paired-end information ran slower than those without paired-end information. Of all the jobs using threads from 16 to 64 on a 32 physical cores system, the speed-up is about 1.8, with about 20 to 25% of the improvement coming from hyper threading.

Figure 7 on page 14 illustrates the performance of two versions of CLC reference mapping on the two System x systems. The algorithm implemented in version 4 beta is completely different from the implementation in version 3, and the result is about 10 times faster than the older version. The performance on x3750 M4 is 10% better than on x3850 X5.

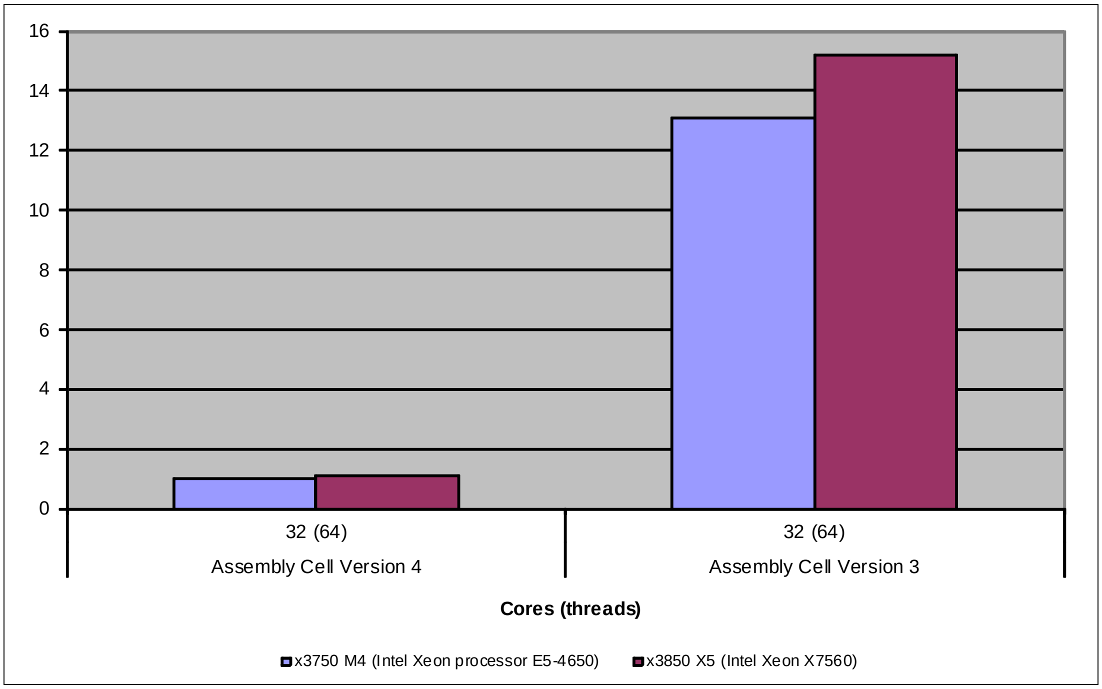

Figure 7 Reference mapping performance results

Most of the NGS applications reached a throughput level of 10^6 bps/s for mapping and alignment, which is in the performance level of version 3 of the CLC Assembly Cell. The version 4 beta of this program is 10 times faster than other mapping and alignment software tools (Hatem and others, 2011: bmi.osu.edu/hpc/papers/Hatem11-Bench.pdf).

Computational chemistry (Gaussian)

Gaussian 09 is the latest version of the Gaussian series of molecular electronic structure programs, used by chemists, chemical engineers, biochemists, physicists, and other scientists worldwide. Gaussian 09 is the most widely used application of its kind and is well known for the wide variety of workloads it can accommodate. The breadth of electronic structure calculations that can be performed by Gaussian 09 makes the application capable of stressing nearly all aspects of computer system performance, including floating point and integer computation, compute-bound performance, memory latency and bandwidth, disk I/O speeds, and both shared and distributed memory parallel capability.

In this study the standard binary distribution of Gaussian 09 Rev.C01 is used to investigate the performance improvement offered by the new IBM System x3750 M4 over a predecessor, the IBM System x3850 X5. Revision C01 is the latest release of Gaussian 09 and includes performance tuning, which takes advantage of the SSE-4 SIMD architecture of both the Intel Xeon processor E5-4650 and Intel Xeon X7560 processors. The test input chosen for the study is a coupled cluster singles-doubles calculation that writes on the order of 100 GB in scratch files—the input is moderately disk I/O intensive. Coupled cluster problems are also known to be both processor and memory-subsystem intensive at various stages throughout the overall calculation. Lastly, recent advances in Gaussian have significantly improved the scaling of coupled cluster calculations. Reasonable single-node scaling was observed in this study. Thus the input chosen for this study represents a good choice for comparing computing systems at the single sever level.

Relative test results (lower is better) at various single-node core counts (8, 16, 32) are shown in Figure 8. Note that time constraints precluded testing at lower core counts.

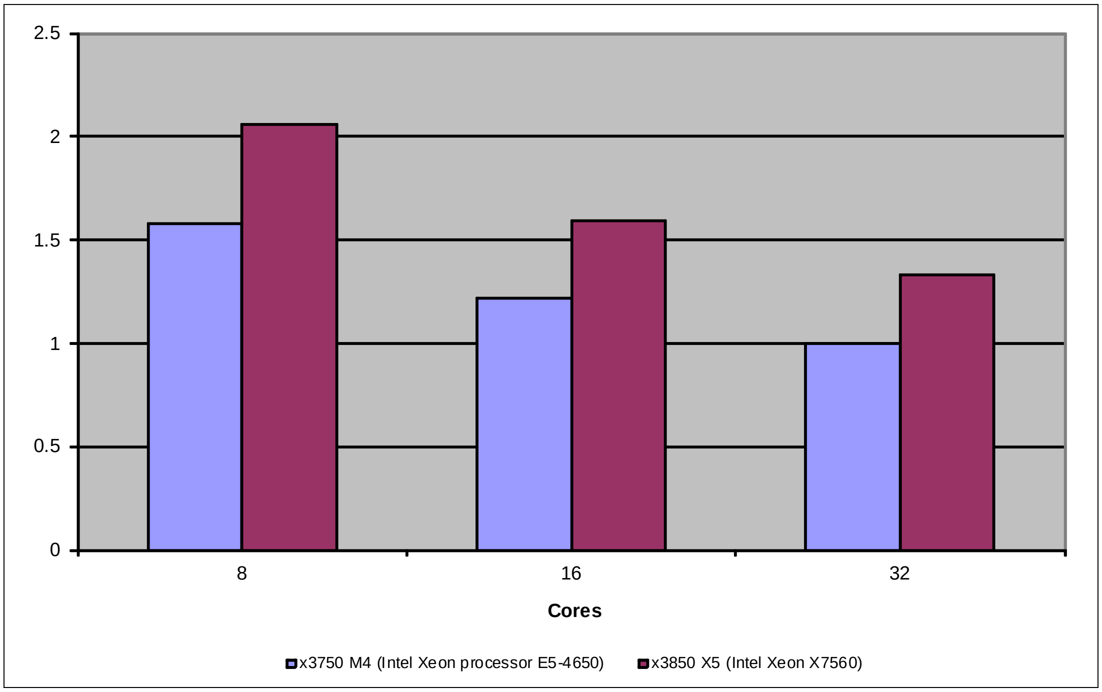

Figure 8 Performance of Gaussian 09 Rev. C.01 on the IBM System x3750 M4

The results show that the performance of the new x3750 M4 is approximately 30% faster across the board than the x3850 for the core counts investigated here. This represents a significant increase in performance on a per-node basis.

Such a performance boost benefits the user in two ways:

- Turnaround times for single jobs are reduced. This means less "waiting for the results" for the user.
- Throughput, that is, the number of test cases that can be run in an extended study, is improved. This allows studies to take less time or allows additional test cases to be run in the same amount of time.

The impact to the client is improved productivity.

As noted at the beginning of this section, Gaussian is well known for the variety of electronic structure calculations it can perform. For example, MP2 and MP4 calculations can also be disk I/O intensive while requiring strong processor and memory-subsystem performance. These types of calculations are more amenable to parallel processing and generally scale to a higher number of cores than the coupled cluster case investigated here. For these calculation types Gaussian 09 can be run across multiple x3750 M4 nodes using LINDA as the distributed-memory parallel API. Lastly, there are a myriad of Gaussian calculation types (Density Functional Theory (DFT), simple SCF, geometry optimization and others) which range in their scalability, memory subsystem requirements, disk I/O usage and need for strong processor performance.

Computer Aided Engineering (FLUENT)

ANSYS Fluent is a computational fluid dynamics (CFD) software solution used to predict fluid flow, heat and mass transfer, chemical reactions, and related phenomena by numerically solving a set of governing mathematical equations (conservation of mass, momentum, energy, and others). ANSYS Fluent, along with other engineering simulation tools from ANSYS, provides engineering teams with product performance data during the conceptual phases of new designs, product development, troubleshooting and redesign.

ANSYS Fluent is both processor- and memory-intensive and hence it is important to consider processor speed as well as scalable memory bandwidth when selecting a compute server to run ANSYS Fluent simulations. Many compute servers may be configured with one processor, two processors, or four processors. IBM System x3750 M4 can contain up to four processors and IBM System x3550 M4 is a two-processor based system. If a single system such as IBM System x3750 M4 can satisfy the performance of a single ANSYS Fluent job with acceptable efficiency, that is preferable to an equivalent cluster-based solution comprising IBM System x3550 M4. IBM System x3750 M4 offers the simplicity of managing a single system image and also avoids the need for adding a complex high-speed network.

The relative results of the FLUENT benchmark for 16 and 32 cores (higher is better) are shown in Figure 9.

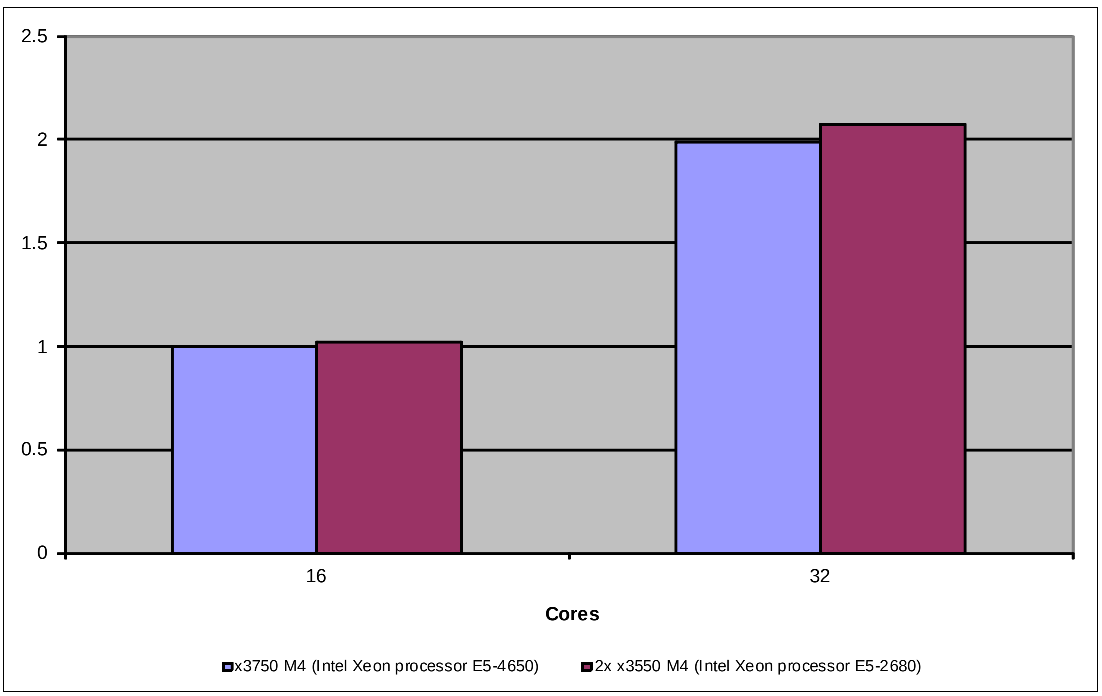

Figure 9 FLUENT benchmark results

The results from the benchmark experiments conducted on IBM System x3750 M4 and equivalent cluster=based solutions clearly establish that the efficiency of IBM System x3750 M4 running ANSYS Fluent is comparable to the IBM System x3550 M4 clusters. For example, performance of x3750 M4 with 32 cores is comparable to a cluster of two x3550 M4 nodes each equipped with 16 cores.

Conclusion

The IBM System x3750 M4 is designed to minimize cost, maximize 2U density, and simplify deployment of HPC workloads, matching performance and capacity of processor, memory, I/O and storage subsystems, lowering overall power consumption, and providing efficient cooling for a dense compute environment.

Key HPC features of the IBM System x3750 M4 include:

- The outstanding RDIMM and LR-DIMM memory performance of the x3750 M4 is achieved by supporting three DIMMs per channel configurations at speeds up to 25% faster than the Intel specification, without sacrificing reliability.
- The use of IBM eXFlash solid-state drives (SSDs) instead of, or along with, traditional spinning drives (HDDs), can improve I/O performance and save energy. An eXFlash can support up to 240,000 I/O operations per second (IOPS) for a total throughput of up to 960,000 IOPS in a single x3750 M4.
- The x3750 M4's 2+2 processor socket design enables pay-as-you-grow processing, memory and I/O scalability with the Intel Xeon processor E5-4600 product family to help lower acquisition cost and manage growth.
- Standard reliability, availability, and serviceability (RAS) features enable easier support and maintenance.

During our benchmark testing, we focused on performance scalability of x3750 M4 itself and in relation to other dual- and quad-socket systems, and observed the following:

- The floating point performance of 4-socket x3750 M4 with Intel Xeon processor E5-4650 is almost equivalent to the performance of two 2-socket dx360 M4 or x3550 M4 with Intel Xeon processor E5-2680.
- Single IBM System x3750 M4 can consolidate an HPC workload from up to five 2-socket dx360 M3 servers with Intel Xeon X5670 processors.
- IBM System x3750 M4 can provide more than three times the technical computing performance compared to the previous generation of x3850 X5 with Intel Xeon X7560 processors.
- The x3750 M4 server can achieve almost linear performance scalability for HPC workloads going from eight to 16 to 32 cores.

While IBM BladeCenter-based and IBM iDataPlex-based HPC cluster deployments are good for large scale implementations, the x3750 M4 is a reasonable choice for small single- or multinode workgroup HPC clusters.

Depending on workload size, x3750 M4 can potentially eliminate the need for multinode HPC workload deployments or significantly reduce the number of nodes required to support specific workloads, therefore providing savings on the external interconnect infrastructure and making the deployment, support and maintenance tasks easier.

Appendix: IBM System x3750 M product details

Table 5 lists the standard specifications of the x3750 M4.

Table 5 Standard specifications of the x3750 M4

Components	Specification
Form factor	2U rack
Processor	Up to four Intel Xeon processor E5-4600 product family processors, each with eight cores (up to 2.7 GHz), six cores (up to 2.9 GHz), or four cores (up to 2.0 GHz). Two processor sockets on the system board and two processors on the processor and memory expansion tray (standard on most models). Two QPI links up to 8.0 GTps each. Up to 1600 MHz memory speed. Up to 20 MB L3 cache per processor.
Chipset	Intel C600 series
Memory	Up to 48 DIMM sockets (12 DIMMs per processor). RDIMMs and LR-DIMMs (Load Reduced DIMMs) are supported, but memory types cannot be intermixed. The memory speed is up to 1600 MHz. There are 24 DIMM sockets on the system board. There are an additional 24 DIMM sockets on the processor and memory expansion tray (standard on most models).
Memory maximums	With RDIMMs: Up to 768 GB with 48x 16 GB RDIMMs and four processors. With LR-DIMMs: Up to 1.5 TB with 48x 32 GB LR-DIMMs and four processors.
Memory protection	ECC, Chipkill (for x4-based memory DIMMs), memory mirroring, and memory sparing.
Disk drive bays	Up to 16 2.5-inch hot-swap SAS/SATA bays or up to 32 1.8-inch hot-swap solid-state drive (SSD) eXFlash bays. Drive bays can be in any combination of four 2.5-inch drives or eight 1.8-inch eXFlash SSD drives.
Maximum internal storage	Up to 14.4 TB with 900 GB 2.5-inch SAS HDDs, up to 16 TB with 1 TB 2.5-inch NL SAS/SATA HDDs, or up to 6.4 TB of SSDs. An intermix of SAS/SATA is supported.
RAID support	RAID 0, 1, 10 with integrated ServeRAID M5110e with LSI SAS2208 6 Gbps RAID on Chip (ROC) controller. Optional upgrades to RAID 5 and 50 are available (zero-cache is 512 MB and battery-backed cache is 512 MB or 1 GB flash-backed cache). There is an optional upgrade to RAID 6 and 60 for a 512 MB or 1 GB cache.
Optical drive bays	There is one bay for an optional multi-burner drive.
Tape drive bays	None internal. Use a supported external tape drive.
Network interfaces	Emulex BE3 controller with two standard integrated Gigabit Ethernet 1000BASE-T ports (RJ-45) and two optional 10 Gb ports through an adapter in a dedicated slot. The 10 GbE options are 10Base-T dual port (copper) or SFP+ dual port (fiber).
PCI Expansion slots	Up to eight slots, five on the system board, three on an optional riser card. The slots are as follows: ► Slot 1: PCIe 3.0 x8; full-height, half-length (optional with riser card, requires CPU 2) ► Slot 2: PCIe 3.0 x8; full-height, half-length (optional with riser card, requires CPU 2) ► Slot 3: PCIe 3.0 x8; full-height, half-length (optional with riser card, requires CPU 2) ► Slot 4: PCIe 3.0 x8; low profile (requires CPU 2) ► Slot 5: PCIe 3.0 x8; low profile (requires CPU 2) ► Slot 6: PCIe 3.0 x8; low profile ► Slot 7: PCIe 3.0 x8; low profile ► Slot 8: PCIe 3.0 x8; low profile Slots 1, 2, and 3 are physically x16 slots.

Components	Specification
Ports	Front: Two USB 2.0 and one DB-15 video on front. Rear: Two USB 2.0, one DB-15 video, one DB-9 serial, one RJ-45 systems management ports, two RJ-45 1 GbE network ports, two optional RJ-45 or SFP+ 10 GbE network ports. Internal: Two internal USB ports (for the embedded hypervisor).
Cooling	IBM Calibrated Vectored Cooling™ with up to six N+N redundant hot swap fans (all six standard); each fan has two rotors.
Power supply	Up to two hot-swap redundant 1400 W AC power supplies (80 PLUS Platinum certification). 900 W power supplies also available through CTO or Special Bid. A second power supply requires that the processor expansion tray or the power interposer card be installed.
Video	Matrox G200eR2 with 16 MB memory integrated into the IMM2. Maximum resolution is 1600x1200 at 75 Hz with 16 M colors.
Hot-swap parts	Hard drives, power supplies, and fans.
Systems management	UEFI, IBM Integrated Management Module II (IMM2), Predictive Failure Analysis, Light Path Diagnostics, Automatic Server Restart, IBM Systems Director and Active Energy Manager, and the IBM ServerGuide. IMM Advanced Upgrade software feature for remote presence are standard with the x3750 M4.
Security features	Power-on password, administrator's password, Trusted Platform Module (TPM).
Operating systems supported	Microsoft Windows Server 2008 R2 and 2008, Red Hat Enterprise Linux 5 and 6, SUSE Linux Enterprise Server 10 and 11, VMware ESX 4.1 and VMware ESXi 4.1 embedded hypervisor, and VMware vSphere 5.
Limited warranty	Three-year customer-replaceable unit and on-site limited warranty with 9x5 next business day (NBD).
Service and support	Optional service upgrades are available through IBM ServicePac® offerings: Four-hour or two-hour response time, eight-hour fix time, one-year or two-year warranty extension, remote technical support for IBM hardware and some IBM and third-party applications.
Dimensions	Height: 86 mm (3.4 in.), width: 445 mm (17.5 in.), depth: 746 mm (29.4 in.)
Weight	Minimum configuration: 25 kg (55 lb.), maximum: 30 kg (65 lb.)

Figure 10 shows the block diagram of the x3750 M4 system.

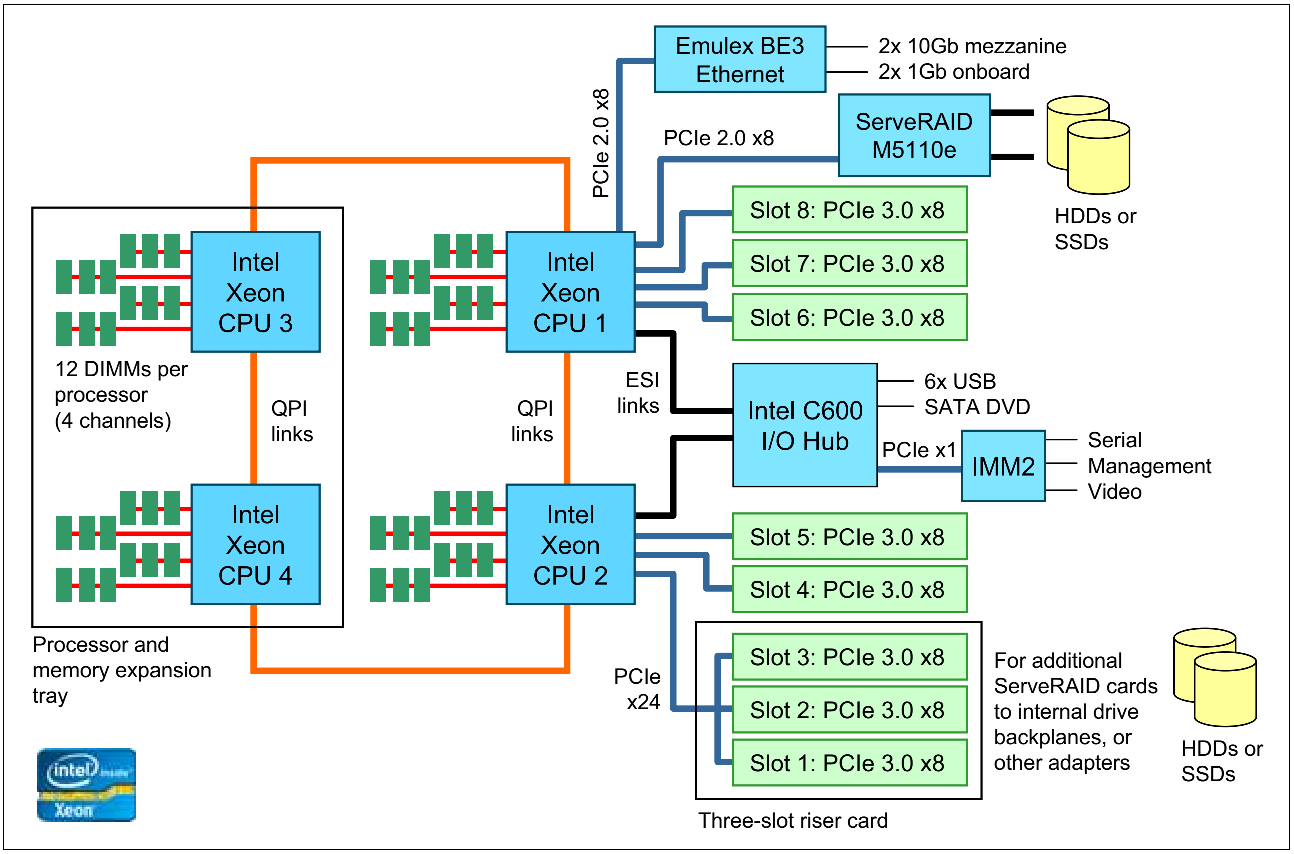

Figure 10 Block diagram of the IBM System x3750 M4

The following technology advances make the x3750 M4 uniquely positioned for HPC workloads:

- Memory speeds that are up to 25% higher than those found in Intel's specification
- IBM eXFlash SSD-based local storage that handles up to 960,000 IOPS
- Pay-as-you-grow 2+2 design approach to scale processor, memory and I/O subsystems
- Dense 2U energy-efficient package
- Proven standard reliability, availability, and serviceability (RAS) features of IBM servers

Memory speeds

One of the key features of the IBM System x3750 M4 is the ability to support memory speeds above the Intel specification. Specifically, in three DIMMs per channel configurations, the following maximum memory speeds are supported:

- 1333 MHz for RDIMMs operating at 1.5 V, which is 25% more than Intel specification establishes
- 1333 MHz for LR-DIMMs operating at either 1.5 V or 1.35 V, which is 25% more than Intel specification establishes

This feature is the result of the IBM design approach, specifically:

- x3750 M4 memory busses are shorter than those required by the Intel specification. The result is a memory bus with less signal loss as the memory data makes its way from the processors to the DIMMs. The shorter bus with less loss results in faster speeds.

- When running at the higher speeds, IBM thoroughly tests each and every DIMM we support. Each DIMM is taken through margin analysis in which voltages, timings and temperatures of the memory bus are varied to ensure that it still meets Intel's specification.
- The processors are kept very centered in relation to memory. By centrally locating the processor within the width of the DIMM, those shorter routes were achieved.

IBM eXFlash

eXFlash technology is a server-based high performance internal storage solution that is based on Solid State Drives (SSDs) and performance-optimized RAID controllers. A single eXFlash unit accommodates up to eight hot-swap SSDs, and can be connected to a single performance-optimized controller.

Each eXFlash unit occupies four 2.5-inch SAS hard disk drive bays. You can install up to four eXFlash units in a single x3750 M4 for up to 32 1.8-inch SSDs (eight SSDs per eXFlash unit).

A single IBM eXFlash unit has the following characteristics:

- Up to eight 1.8-inch hot-swap front-accessible SSDs
- Up to 240,000 random read IOPS
- Up to 2 GB/sec of sustained read throughput
- Up to 1.6 TB of available storage space with IBM 200 GB 1.8-inch eMLC SSDs or up to 400 GB with IBM 50 GB 1.8-inch eMLC SSDs

In theory, the random I/O performance of a single eXFlash unit is equivalent to the storage system consisting of about 800 traditional spinning HDDs. Besides HDDs themselves, building such a massive I/O-intensive high-performance storage system requires external deployment with many additional infrastructure components, including host bus adapters (HBAs), switches, storage controllers, disk expansion enclosures, and cables. Consequently, this leads to more capital expenses, floor space, electrical power requirements, and operation and support costs. Because eXFlash is an internal server storage, it does not require all these components, and helps to eliminate additional expenses and environmental requirements.

In summary, the IBM eXFlash solution provides the following benefits:

- Significantly lower implementation cost (up to 97% lower) of high performance I/O-intensive storage systems with best price/IOPS performance ratio
- Significantly higher performance (up to 30 times or more) of I/O-intensive applications such as databases and business analytics with up to nine times shorter response time
- Significant savings in power and cooling (up to 90%) with high performance per watt ratio
- Significant savings in floor space (up to 30 times less) with extreme performance per U ratio
- Simplified management and maintenance with internal server-based configurations (no external power and information infrastructure needed).

In addition to its superior performance, eXFlash offers superior uptime with three times the reliability of mechanical disk drives. SSDs have no moving parts to fail. They use Enterprise Wear-Leveling to extend their use even longer. All operating systems that are listed in IBM ServerProven® for the x3750 M4 are supported for use with eXFlash.

Pay-as-you-grow scalability

The IBM System x3750 M4 server design utilizes a 2+2 scalability approach. You can start from a dual-socket configuration with 24 DIMM slots and five PCIe 3.0 slots. When needed, you can upgrade the existing configuration by adding the processor and memory expansion tray that provides two additional processor sockets and 24 DIMM slots for a total of four processors and up to 1.5 TB of memory with 32 GB LR-DIMMs that all run at up to 1333 MHz. Optionally, you can also scale the I/O subsystem by adding three more PCIe 3.0 slots.

From a performance standpoint, such an approach provides good scalability from one to two to four processors, as we describe in “HPC benchmarks” on page 5. Such vertical scaling can help eliminate or significantly reduce a costly high-speed cluster interconnect infrastructure for the majority of mid-sized HPC workloads.

Other performance and scalability features include:

- The Intel Xeon processor E5-4600 product family improves productivity by offering superior system performance with 8-core processors and up to 2.7 GHz core speeds (8-core processors), up to 20 MB of L3 cache, and up to two 8 GTps QPI interconnect links.
- Up to four processors, 32 cores, and 64 threads with Intel Hyper-Threading Technology maximize the concurrent execution of multithreaded applications.
- Intelligent and adaptive system performance with Intel Turbo Boost Technology 2.0 allows processor cores to run at maximum speeds during peak workloads by temporarily going beyond processor TDP.
- Intel Advanced Vector Extensions (AVX) can improve floating-point performance for compute-intensive technical and scientific applications.
- Up to 16 HDDs or 32 eXFlash SSDs, together with an optical drive at the same time, provide a flexible and scalable all-in-one platform to meet your increasing demands.
- The server has two integrated Gigabit Ethernet ports and two optional 10 Gb Ethernet ports that do not consume PCIe slots.
- The server offers PCI Express 3.0 I/O expansion capabilities that improve the theoretical maximum bandwidth by almost 100% (8 GTps per link using 128b/130b encoding) compared to the previous generation of PCI Express 2.0 (5 GTps per link using 8b/10b encoding).
- With Intel Integrated I/O Technology, the PCI Express 3.0 controller is integrated into the Intel Xeon processor E5 family. This integration reduces I/O latency and increases overall system performance.

Energy efficiency

The x3750 M4 offers the following energy-efficiency features to save energy, reduce operational costs, increase energy availability, and contribute to a green environment:

- Energy-efficient planar components help lower operational costs.
- Highly efficient 900 W and 1400 W AC power supplies with 80 PLUS Platinum certification at high voltage AC.
- The Intel Xeon processor E5-4600 product family offers balanced performance and energy use with built-in Intel TurboBoost 2.0 and Intel Intelligent Power Capability technologies.
- Low-voltage 1.35 V DDR3 memory DIMMs consume 19% less energy than 1.5 V DDR3 RDIMMs.

- Solid-state drives (SSDs) consume as much as 80% less power than traditional spinning 2.5-inch HDDs.
- The server uses hexagonal ventilation holes, which is a part of IBM Calibrated Vectored Cooling technology. Hexagonal holes can be grouped more densely than round holes, providing more efficient airflow through the system.
- IBM Systems Director Active Energy Manager™ provides advanced data center power notification and management to help achieve lower heat output and reduced cooling needs.

RAS features

The x3750 M4 provides many features to simplify serviceability and increase system uptime:

- The server offers Chipkill, memory mirroring and memory rank sparing for redundancy in the event of a memory failure.
- The server provides restart recovery for any failed processor. In the event of a failure of processor 1, the server connects the south bridge to processor 2 for reboot.
- The server has up to two redundant hot-swap power supplies and six hot-swap dual-rotor N+N redundant fans to provide availability for business-critical applications.
- The power source-independent light path diagnostics panel and individual light path LEDs lead the technician to failed (or failing) components, which simplifies servicing, speeds up problem resolution, and helps improve system availability.
- Predictive Failure Analysis (PFA) detects when system components operate outside of standard thresholds and generates proactive alerts in advance of a possible failure, therefore increasing uptime. These components support PFA:
  - Memory
  - SAS/SATA HDDs
  - Fans
  - VRDs
  - Power supplies
- Solid-state drives (SSDs) offer more reliability than traditional mechanical HDDs for greater uptime.
- The built-in Integrated Management Module Version II (IMM2) continuously monitors system parameters, triggers alerts, and performs recovery actions in case of failures to minimize downtime.
- Built-in diagnostics, using Dynamic Systems Analysis (DSA) Preboot, speed up troubleshooting tasks to reduce service time.
- Three-year customer-replaceable unit and on-site limited warranty, 9x5 next business day. Optional service upgrades are available.

For more information, refer to *IBM System x3750 M4 Implementation Guide*, REDP-4874.

The team who wrote this paper

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization, Raleigh Center.

Ilya Krutov is a Project Leader at the ITSO Center in Raleigh and has been with IBM since 1998. Before joining the ITSO, Ilya served in IBM as a Run Rate Team Leader, Portfolio Manager, Brand Manager, Technical Sales Specialist, and Certified Instructor. Ilya has expertise in IBM System x and BladeCenter products, server operating systems, and networking solutions. He has a bachelor degree in computer engineering from the Moscow Engineering and Physics Institute.

Raj Panda is the Team Leader of the HPC Performance Team in Systems and Technology Group and has been with IBM since 2001. Raj works in the area of parallel application tuning and optimization with expertise in application performance projection on future generation parallel computers. He has a Ph.D. in Mechanical and Aerospace Engineering from Princeton University.

Arthur Ban is a Performance Lab Leader at the STG Power R&D Lab in Austin and has been with IBM since 1987. Before joining the STG, Arthur served in IBM as an OfficeVision/MVS development Project Leader, OfficeVision/CICS performance analyst, and OS/2 system performance analyst. Arthur has development and performance area expertise in numerous IBM systems ranging from PC to mainframe and all Power system products. He has a Master of Science degree in computer science from the University of Texas at Arlington.

Ruzhu Chen is an IBM Certified Expert IT Specialist in the IBM HPC life sciences solution team. He leads technical support for computational chemistry and NGS applications such as NWChem, AMBER, NAMD, CLC Assembly Cell, and BLAST. He provides application enablement, tuning, optimization and benchmarks for IBM world-wide partners and clients. He has more than 10 years of expertise in HPC application tuning and optimization, and has engaged in numeral application solutions enablements on IBM Power and x86_64 architectures. Ruzhu holds B.S. and M.S degrees in Biochemistry from the University of Sciences and Technology of China, and an M.S. in Computer Science and Ph.D. in Molecular Biology from the University of Oklahoma.

Zaphiris Christidis is a Senior Certified IT Architect in STG Systems Software Group. He has been in IBM since 1987. Before joining STG in 2009, Zaphiris was a Research Staff Member at the TJ Watson Research Lab, in the Deep Computing and Parallel Systems departments. In 2002 Zaphiris joined STG in the Greater China Group as an International assignee, working on High Performance Computing and Environmental Applications. Zaphiris has expertise in designing weather system solutions, benchmarking and Weather code optimization on system P and X products. Zaphiris has a Bachelor's degree in Mathematics from the University of Ioannina, Greece, an MS in Climatology from the City College of New York, and a Ph.D. in Atmospheric Sciences and Engineering from the University of Michigan.

Bradley Elkin is a Senior Development Engineer with the Server Hardware and HPC Performance Group and has been with IBM since 2001. His areas of expertise include hardware performance analysis and the performance optimization of technical computing workloads. He holds a Ph.D. in Chemical Engineering from the University of Pennsylvania and has 25 years of experience in high performance computing. He has written several articles for the IBM eServer™ Development Domain.

Tony Pirraglia is a senior engineer in the ISV Solution Enablement organization. He is based in Poughkeepsie, NY and is currently involved with ISV application benchmarks, interfacing with the IBM Blue Gene® team regarding ISV support and application benchmarking, and representing the ISV organization on selected Product Development Teams. Tony has been involved with Scientific and Technical computing on IBM platforms in one form or another since 1988, when he first joined IBM. He received his Doctorate in Chemical Engineering from Columbia University.

Hari Reddy is a Senior Technical Consultant with Global ISV Technical Enablement within the IBM Systems and Technology Group. Hari joined IBM in 1990. His areas of expertise include porting and tuning large scale parallel applications. He holds an M.S. degree in Operations Research and a Ph.D. in Computer Science.

Thanks to the following people for their contributions to this project:

Tamikia Barrow, Mary Comianos, Shari Deiana, Linda Robinson, David Watts,
International Technical Support Organization, Raleigh Center

Kevin Barnes, Alfred Schwab, editor, Cheryl Gera, Erica Wazewski
International Technical Support Organization, Poughkeepsie Center

Randall Lundin
IBM

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new IBM Redbooks® publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document REDP-4890-00 was created or updated on November 28, 2012.

Send us your comments in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:
  ibm.com/redbooks
- Send your comments in an email to:
  redbooks@us.ibm.com
- Mail your comments to:
  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400 U.S.A.

Redpaper ™

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

BladeCenter®
Blue Gene®
Calibrated Vectored Cooling™
eServer™
GPFS™
IBM Systems Director Active Energy Manager™
IBM®
iDataPlex®
Redbooks®
Redpaper™
Redbooks (logo) ®
ServerProven®
ServicePac®
System x®

The following terms are trademarks of other companies:

Intel Xeon, Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.